

CLAIMS

What is claimed is:

1. 1. A method for updating a shared data element group while preserving group integrity on behalf of one or more readers that are concurrently referencing group data elements without using locks or atomic instructions, comprising:
 - 4 generating a new group data element;
 - 5 assigning a generation number to said new data element that allows a reader of said
 - 6 data element group to determine whether said new data element is a correct version for said
 - 7 reader;
 - 8 if a prior version of said new data element exists, establishing a version link between
 - 9 said new data element and said prior version;
 - 10 linking said new data element into said data element group so that it is reachable by
 - 11 readers;
 - 12 updating a global generation number associated with said data element group; and
 - 13 if a prior version of said new data element exists, freeing said prior version following
 - 14 a grace period.
1. 2. A method in accordance with claim 1 wherein said method is used to replace a group data element and said new data element is generated by copying said data element to be replaced.
1. 3. A method in accordance with claim 1 wherein said method is used to delete a group data element and said new data element is generated by copying said data element to be deleted and setting a deletion flag in said new data element.
1. 4. A method in accordance with claim 1 wherein said method is used to insert a new

2 group data element and said new data element has no prior versions.

1 5. A method in accordance with claim 1 wherein said method further includes generating
2 a pointer-forwarding entity that points to said new data element, said pointer forwarding
3 entity maintaining said version link on behalf of said new data element and further being used
4 to link said new data element into said data element group.

1 6. A method for updating a shared data element group while preserving group integrity
2 on behalf of one or more readers that are concurrently referencing group data elements
3 without using locks or atomic instructions, comprising:

4 generating a pointer-forwarding entity that points to a data element in said data
5 element group;

6 assigning a generation number to said pointer-forwarding entity that allows a reader of
7 said data element group to determine whether said pointer-forwarding entity is a correct
8 version for said reader;

9 if there is a prior version of said pointer-forwarding entity, establishing a version link
10 between said pointer-forwarding entity and said prior version;

11 linking said pointer-forwarding entity into said data element group so that said data
12 element pointed to by said pointer-forwarding entity is reachable by readers through said
13 pointer-forwarding entity;

14 updating a global generation number associated with said data element group; and
15 if a prior version of said pointer-forwarding entity exists, freeing said prior version
16 following a grace period.

1 7. A method for performing a search of a shared data element group that may be
2 undergoing modification in accordance with the method of claim 1, comprising:

3 assigning a current global generation number to said search;
4 when referencing a data element in said data element group, determining whether said
5 referenced data element is a correct version by comparing a generation number assigned to
6 said referenced data element with said search generation number; and
7 searching for a correct version of said referenced data element as necessary.

1 8. A method in accordance with Claim 7 wherein, if said data element generation number
2 is equal to said search generation number, said referenced data element is accepted for reading
3 as a correct version.

1 9. A method in accordance with Claim 7 wherein, if said data element generation number
2 is less than said search generation number, a search is made for a later version of said
3 referenced data element, and wherein said referenced data element is used if a later version is
4 not found.

1 10. A method in accordance with Claim 7 wherein, if said data element generation number
2 is greater than said search generation number, a search is made for a prior version of said
3 referenced data element, and wherein said referenced data element is deemed to be a new
4 insertion if there is no prior version.

1 11. A data processing system having one or more central processing units, a memory and
2 a communication pathway between the one or more central processing units and the memory,
3 said system being adapted to update a shared data element group in said memory while
4 preserving group integrity on behalf of one or more readers that are concurrently referencing
5 group data elements without using locks or atomic instructions, and comprising:
6 means for generating a new group data element;
7 means for assigning a generation number to said new data element that allows a reader

8 of said data element group to determine whether said new data element is a correct version for
9 said reader;

10 means for establishing, if a prior version of said new data element exists, a version
11 link between said new data element and said prior version;

12 means for linking said new data element into said data element group so that it is
13 reachable by readers;

14 means for updating a global generation number associated with said data element
15 group; and

16 means for freeing, if a prior version of said new data element exists, said prior version
17 following a grace period.

1 12. A system in accordance with claim 11 wherein said system is adapted to replace a
2 group data element and to generate said new data element by copying said data element to be
3 replaced.

1 13. A system in accordance with claim 11 wherein said system is adapted to delete a
2 group data element and to generate said new data element by copying said data element to be
3 deleted and setting a deletion flag in said new data element.

1 14. A system in accordance with claim 11 wherein said system is adapted to insert a new
2 group data element such that said new data element has no prior versions.

1 15. A system in accordance with claim 11 wherein said system further includes means for
2 generating a pointer-forwarding entity that points to said new data element, said pointer
3 forwarding entity maintaining said version link on behalf of said new data element and further
4 being used to link said new data element into said data element group.

1 16. A data processing system having one or more central processing units, a memory and
2 a communication pathway between the one or more central processing units and the memory,
3 said system being adapted to update a shared data element group in said memory while
4 preserving group integrity on behalf of one or more readers that are concurrently referencing
5 group data elements without using locks or atomic instructions, and comprising:
6 means for generating a pointer-forwarding entity that points to a data element in said
7 data element group;
8 means for assigning a generation number to said pointer-forwarding entity that allows
9 a reader of said data element group to determine whether said pointer-forwarding entity is a
10 correct version for said reader;
11 means for establishing, if there is a prior version of said pointer-forwarding entity, a
12 version link between said pointer-forwarding entity and said prior version;
13 means for linking said pointer-forwarding entity into said data element group so that
14 said data element pointed to by said pointer-forwarding entity is reachable by readers through
15 said pointer-forwarding entity;
16 means for updating a global generation number associated with said data element
17 group; and
18 means for freeing, if a prior version of said pointer-forwarding entity exists, said prior
19 version following a grace period.

1 17. A data processing system having one or more central processing units, a memory and
2 a communication pathway between the one or more central processing units and the memory,
3 said system being adapted to perform a search of a shared data element group that may be
4 undergoing modification by the system the system of claim 11, and comprising:

5 means for assigning a current global generation number to said search;
6 means for determining, when referencing a data element in said data element group,
7 whether said referenced data element is a correct version by comparing a generation number
8 assigned to said referenced data element with said search generation number; and
9 means for searching for a correct version of said referenced data element as necessary.

1 18. A system in accordance with Claim 17 wherein said system is adapted to accept said
2 referenced data element for reading as a correct version if said data element generation
3 number is equal to said search generation number.

1 19. A system in accordance with Claim 17 wherein said system is adapted to search for a
2 later version of said referenced data element if said data element generation number is less
3 than said search generation number, and to use said referenced data element if a later version
4 is not found.

1 20. A system in accordance with Claim 17 wherein said system is adapted to search for a
2 prior version of said referenced data element if said data element generation number is greater
3 than said search generation number, and to deem said referenced data element to be a new
4 insertion if there is no prior version.

1 21. A computer program product for updating a shared data element group while
2 preserving group integrity on behalf of one or more readers that are concurrently referencing
3 group data elements without using locks or atomic instructions, comprising:

4 one or more data storage media;
5 means recorded on said data storage media for programming a data processing
6 platform to operate as by:
7 generating a new group data element;

8 assigning a generation number to said new data element that allows a reader of said
9 data element group to determine whether said new data element is a correct version for said
10 reader;

11 if a prior version of said new data element exists, establishing a version link between
12 said new data element and said prior version;

13 linking said new data element into said data element group so that it is reachable by
14 readers;

15 updating a global generation number associated with said data element group; and

16 if a prior version of said new data element exists, freeing said prior version following
17 a grace period.

1 22. A computer program product in accordance with claim 21 wherein said product is
2 adapted to replace a group data element and to generate said new data element by copying
3 said data element to be replaced.

1 23. A computer program product in accordance with claim 21 wherein said product is
2 adapted to delete a group data element and to generate said new data element by copying said
3 data element to be deleted and setting a deletion flag in said new data element.

1 24. A computer program product in accordance with claim 21 wherein said product is
2 adapted to insert a new group data element such that said new data element has no prior
3 versions.

1 25. A computer program product in accordance with claim 21 wherein said means
2 recorded on said data storage media are further adapted to generate a pointer-forwarding
3 entity that points to said new data element, said pointer forwarding entity maintaining said

4 version link on behalf of said new data element and further being used to link said new data
5 element into said data element group.

1 26. A computer program product for updating a shared data element group while
2 preserving group integrity on behalf of one or more readers that are concurrently referencing
3 group data elements without using locks or atomic instructions, comprising:

4 one or more data storage media;

5 means recorded on said data storage media for programming a data processing
6 platform to operate as by:

7 generating a pointer-forwarding entity that points to a data element in said data
8 element group;

9 assigning a generation number to said pointer-forwarding entity that allows a reader of
10 said data element group to determine whether said pointer-forwarding entity is valid for said
11 reader;

12 if there is a prior version of said pointer-forwarding entity, establishing a version link
13 between said pointer-forwarding entity and said prior version;

14 linking said pointer-forwarding entity into said data element group so that said data
15 element pointed to by said pointer-forwarding entity is reachable by readers through said
16 pointer-forwarding entity;

17 updating a global generation number associated with said data element group; and
18 if a prior version of said pointer-forwarding entity exists, freeing said prior version
19 following a grace period.

1 27. A computer program product for performing a search of a shared data element group
2 that may be undergoing modification in accordance with the method of claim 1, comprising:

3 one or more data storage media;

4 means recorded on said data storage media for programming a data processing

5 platform to operate as by:

6 assigning a current global generation number to said search;

7 when referencing a data element in said data element group, determining whether said

8 referenced data element is a correct version by comparing a generation number assigned to

9 said referenced data element with said search generation number; and

10 searching for a correct version of said referenced data element as necessary.

1 28. A computer program product in accordance with Claim 27 wherein, if said data

2 element generation number is equal to said search generation number, said referenced data

3 element is accepted for reading as a correct version.

1 29. A computer program product in accordance with Claim 27 wherein, if said data

2 element generation number is less than said search generation number, a search is made for a

3 later version of said referenced data element, and wherein said referenced data element is used

4 if a later version is not found.

1 30. A computer program product in accordance with Claim 27 wherein, if said data

2 element generation number is greater than said search generation number, a search is made for

3 a prior version of said referenced data element, and wherein said referenced data element is

4 deemed to be a new insertion if there is no prior version.

31. A computer program product for managing a shared data element group so as to allow updates thereof while preserving group integrity on behalf of one or more readers that are concurrently referencing group data elements without using locks or atomic instructions, comprising:

5 one or more data storage media;

means recorded on said data storage media for programming a data processing platform to operate as by:

performing a first-phase update operation that preserves a consistent pre-update view of said data element group on behalf of pre-update readers and a consistent post-update view
10 of the data element group on behalf of post-update readers;

providing means by which readers can locate all data elements of said data element group that belong to each of said pre-update and post-update views as readers search said data element group;

15 performing one or more read operations following said first-phase update operation in which one or more readers search said data element group with each reader referencing only data elements belonging to one of said pre-update and post-update views; and

performing a second-phase update operation following a grace period that frees said pre-update view of said data element group.